The Performance of the NAS HSPs in 4Q93

Robert J. Bergeron

Report RND-94-012 July 1994

NAS Systems Development Branch
NAS Systems Division
NASA Ames Research Center
Mail Stop 258-6
Moffett Field, CA 94035-1000

Abstract

During 4Q93, the NAS C-90 delivered an average throughput of 3.626 GFLOPS while the NAS Y-MP averaged 0.678 GFLOPS. The C-90 rate was 5% greater than the previous two quarters. Reasons for the increased performance included a vector length which exceeded the 3Q value by 13% and a vector operation fraction which exceeded the 3Q value by 1%. Despite the increase in vector length, insufficient vector lengths continue to hinder HSP performance. To improve HSP performance, NAS should encourage the HSP users to modify their codes to increase program vector length.

1.0 Introduction

The introduction of the C-90 in March 1993 motivated the daily monitoring of the hardware performance of the NAS High Speed Processors (HSPs). This paper, covering the 4th quarter of 1993, is the third report in the series.

The C-90 Hardware Performance Monitor (HPM)continuously delivers a full 32-counter record for the workload [1]. The Y-MP HPM can report only a single group (8-counter) during any one period. Since NAS monitors the Group 0 performance of every program in the workload, only the Group 0 counters will be reported here.

NAS records the daily average values of all HPM counters. In 4Q93, maintenance activities on the C-90 in late December permitted 88 days of user service. This report provides tables of counter values representing the average, maximum, and minimum values from the 88 daily reports in this quarter. The removal of the Y-MP from the NAS facility in late October permitted 30 days of user service in this quarter and the report provides tables of counter values representing the average, maximum, and minimum values from the 30 daily reports in 4Q93.

The tables provide performance rate data per CPU for the actual time the CPU spent executing the user programs. System throughput is based on wall clock time and total number of CPUs. A complete explanation of all counter data occurs in [2].

To provide a feel for the daily variation in each of the counters, the report also provides the standard deviation (STD) and coefficient of variation (COV). The coefficient of variation is the ratio of the standard deviation of a quantity divided by its average value.

2.0 C-90 Counter Data

The NAS C-90 (SN 4012) CPUs have a clock period (CP) of 4.167 nanoseconds and a peak speed of 960 MFLOPS. This architecture has 128-element vector registers and double-width 64-element functional units.

For clarity, the report divides the 32 C-90 counters into 4 functional groups: global counters, instruction holds, instruction issues, and vector operations.

Global Counters:

Table 1 provides counter data giving a total counts for instructions, operations and references. The unit "M/sec" denotes "Million per sec" and the unit "avg/ref" denotes "average (conflict) per reference". The term "reference" denotes a single Cray word (8-byte) data transfer.

Table 1: NAS C-90 4Q93 Daily Average HPM Measurements- Global Counters

Measurement	Unit	Avg	STD	COV	Min	Max
CPU time	sec	62534.	15475.	0.247	17782.	82962.
Instruction Issue	M/sec	56.651	2.497	0.044	48.437	61.971
Average clock periods/inst	----	4.245	0.196	0.046	3.872	4.954
CP holding issue	Percent	67.629	1.531	0.023	64.286	72.684
Instruction buffer fetches	M/sec	0.272	0.032	0.119	0.184	0.355
Floating Pt. Ops per CPU	M/sec	255.432	17.912	0.070	214.533	296.727
Vector Floating Pt. Ops	M/sec	250.969	18.286	0.073	209.194	294.100
CPU memory references	M/sec	247.558	17.874	0.072	203.403	311.754
CPU memory conflicts	avg/ref	0.289	0.037	0.129	0.221	0.388
VEC memory references	M/sec	242.005	18.322	0.076	196.660	308.009
B/T memory references	M/sec	1.357	0.250	0.184	0.811	2.069
I/O memory references	M/sec	1.660	0.724	0.436	0.496	3.490
I/O memory conflicts	avg/ref	0.301	0.023	0.078	0.263	0.368

The large variation in the CPU time measurement reflects the requirement that the HPM data represent a continuous interval. Occasionally, persistent hardware and/or software problems may require several shutdowns during the 24-hour measurement period. The CPU time reported in the table for such days is the longest continuous period without a shutdown.

During 4Q93, each of the C-90's 16 CPUs performed at an average rate of 255 MFLOPS, an increase of 4.7% over the previous quarter. The rate of

instruction issue decreased by 5.4% and the average number CPs per instruction increased by 5.6%. Since more work was done with fewer instructions and since the instructions took longer to complete, these two measurements suggest an increased 4Q workload vector fraction. Although the average CPU rate is well below the single CPU maximum of 960 MFLOPS, many NAS applications display performance rates exceeding 500 MFLOPS.

The CPU MFLOP rate exceeds the CPU memory reference rate by a factor of 1.03, indicating that the average floating point operation must be reusing data in the registers to avoid memory accesses. The CPU memory reference rate was 6.8% higher than the previous quarter's rate and the measurements indicate a 9.4% increase in the number of conflicts per memory references. Measurements discussed in the next section indicate that memory access is not yet a bottleneck for this workload.

The I/O memory reference rate of 1.66 Mwords/sec per CPU represents a 47% decrease in I/O related data transfers. HSP monitoring indicates that peaks in the I/O transfer rate are strongly related to the restores required by failures of the various system disks. Examination of system logs indicates fewer restores in 4Q93. Monitoring of the I/O to the disks in December indicate a maximum of 5.0 Mwords/sec implying that about 90% of the C-90 I/O targets the SSD. Although the average I/O rate is well below the single CPU maximum of 239 Mwords/sec, several NAS applications sustain data transfer rates of 200 Mwords/sec.

The I/O rate measurements display a large COV relative to the performance rate measurements.This variance reflects the differing input/output requirements of NAS users. The Cray timesharing architecture decouples the I/O rate from the MFLOP rate because the data transfer occurs when the user program has given up control of the CPU to another program. The second program can maintain the CPU MFLOP rate while the I/O from the first program proceeds. If the transfer is efficient and the two programs have similar performance characteristics, measurements should show the MFLOP rate relatively constant while the I/O rate fluctuates according to user needs. The C-90 measurements substantiate this claim.

Instruction Holds:

Instructions are fetched from the instruction buffer by the instruction processor. If any of the resources required to execute the instruction are reserved, the instruction issue logic prevents the instruction from issuing. The HPM records all CPs for which the instruction holds issue and the table presents these as the percent of total CPU time. Since there may be more than one resource reservation preventing an instruction issue, the sum of the percentages in this group can exceed 100%.

Table 2: NAS C-90 4Q93 Daily Average HPM Measurements- Instruction Holds

Measurement	Unit	Avg	STD	COV	Min	Max
Waiting on A-registers	% CPU	4.842	0.363	0.075	4.101	6.350
Waiting on S-registers	% CPU	9.531	1.201	0.126	6.296	12.317
Waiting on V-registers	% CPU	22.261	1.676	0.075	19.375	28.249
Waiting on B/T-registers	% CPU	1.315	0.198	0.151	0.836	2.017
Waiting on F'nctnal Units	% CPU	24.384	1.518	0.062	21.732	28.682
Waiting on Shared Regs	% CPU	0.386	0.250	0.647	0.014	1.089
Waiting on Memory Ports	% CPU	16.725	1.577	0.094	13.466	21.141
Waiting on Miscellaneous	% CPU	2.326	0.093	0.040	2.045	2.547

The major reason for instruction issue delays are busy vector registers and busy vector functional units; the instruction processor will not issue an instruction until operations in these units have completed. Calculations derived from counter data (Table 5) have shown that other operations were in progress during these delays.

The approximately equal delays in vector registers and vector functional units indicates efficient register use and overlapping of vector functional units.

A CPU memory port accesses a section which accesses a memory bank. Memory references can lead to two kinds of delay in the Y-MP/C-90 architecture. A memory instruction hold occurs, for example, when a register is reserved by another instruction or a memory port is busy. A memory conflict occurs when a needed bank is busy. A user program executing on a single CPU can encounter conflicts when it continuously references the same bank. A workload can encounter conflicts when several CPUs simultaneously reference the same bank.

For the workload vector length of 68, the C-90 can store data to memory at a rate of 1.28 CP/word. Data from Table 1 indicates that each memory reference on the average experiences a memory contention delay of 0.289 CP. Table 2 data indicate that memory resources prevent the CPU from issuing an instruction about 17% of the time and converting this delay to a per reference basis yields a memory delay of 0.162 CP.

Total memory delay is 0.289 + 0.162, or about 0.44 CP/reference and this delay is a fraction of the 1.28 CP minimum required for a C-90 workload vector memory reference.

The following figure shows 1993 C-90 performance as a function of memory delay and indicates a slight decrease in memory delay as performance increases and an average delay of about 0.45 CP, the same as the average for 4Q93.

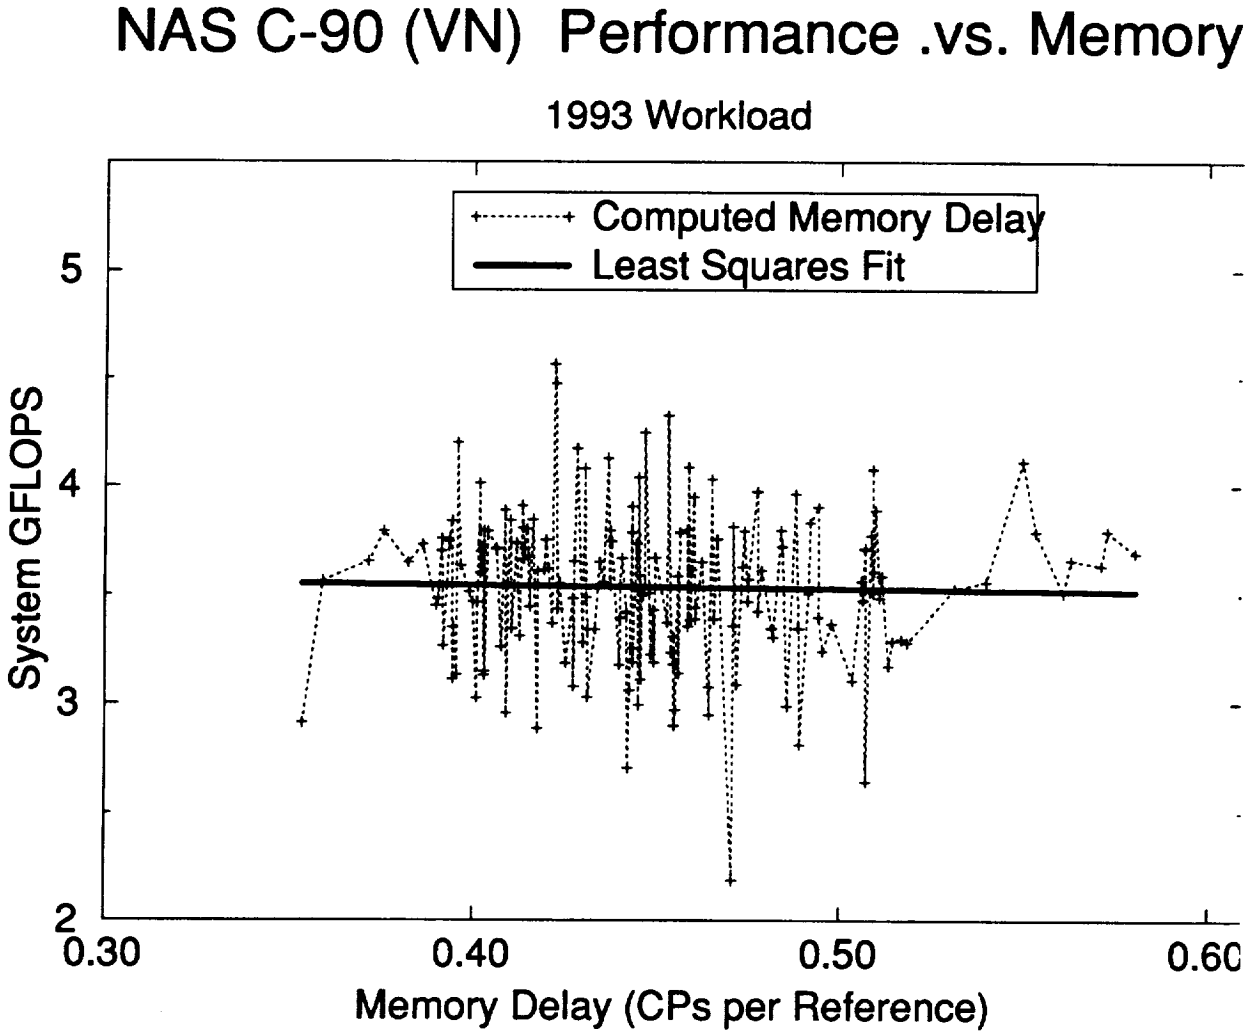

Instruction Issues:

The unit "M/sec" denotes "Millions of instructions per second". Instructions produce the operations which constitute the actual workload tasks.

Table 3: NAS 4Q93 C-90 Daily Average HPM Measurements- Instruction Issues

Measurement	Unit	Avg	STD	COV	Min	Max
(000-004)Special	M/sec	1.104	0.090	0.082	0.892	1.393
(005-017)Branch	M/sec	2.497	0.229	0.092	1.857	3.101
(02x,030-033)A Register	M/sec	23.789	1.183	0.050	21.208	27.629
(034-037)B/T Memory	M/sec	0.158	0.034	0.213	0.091	0.282
(040-043,071-077)S Register	M/sec	7.700	0.836	0.109	5.503	9.427
(044-061)Scalar Integer	M/sec	4.612	0.557	0.121	3.224	6.077
(062-070)Scalar Floating Pt.	M/sec	4.464	0.692	0.155	2.627	6.041
(10x-13x)Scalar Memory	M/sec	4.196	0.500	0.119	2.860	5.137
(140-177)All Vector	M/sec	8.131	0.386	0.047	7.176	9.121

A-register instructions comprise about 42% of the scalar instructions issued. These instructions compute memory addresses and indexes for memory, loop control, and I/O.

The C-90 percentage of A-register instructions exceeds the 32% reported for the NAS Y-MP workload and this increase appears due to the additional address overhead of the multiple pipes. All CPUs of the C-90/Y-MP architecture have two pipes, one consisting of an add functional unit and the other consisting of a multiply functional unit. In both Y-MP and C-90 platforms, the number of segments in the functional unit is 64. However, the C-90 functional unit design provides two 64-element columns and these double-width functional units require some additional A-register operations.

Scalar instructions constitute about 44% of NAS workload instructions and this value represents a larger percentage than 3Q93.

Vector instructions are only 14% of the total instructions, but vector operations represent about 92% of the workload operations (Table 5). A single vector instruction can produce many vector operations.

Vector Operations:

All of the vector operations shown in Table 4 are produced by vector instructions. Table 3 shows that the rate of instruction issue for all vector instructions was 8.131 million per second
The vector operation rate for 4Q93, which is the sum of the column 3 values in the first 8 rows of Table 4, was 549 million per second.

Table 4: NAS C-90 4Q93 Daily Average HPM Measurements- Vector Operations

Measurement	Unit	Avg	STD	COV	Min	Max
Vector Logical	M/sec	31.003	4.222	0.136	20.629	45.701
Vector Shift/Pop/LZ	M/sec	8.072	0.777	0.096	6.592	12.043
Vector Integer Add	M/sec	16.526	1.819	0.110	12.664	23.022
Vector Floating Pt. Multiply	M/sec	127.483	8.913	0.070	104.951	146.882
Vector Floating Pt. Add	M/sec	116.301	9.608	0.083	97.468	140.817
Vector Floating Reciprocal	M/sec	7.185	0.818	0.114	5.326	9.335
Vector Memory Read	M/sec	168.457	12.412	0.074	137.195	207.259
Vector Memory Write	M/sec	73.548	6.299	0.086	59.464	100.750
Average Vector Length	----	67.584	4.997	0.074	55.590	79.270

The ratio of total vector operations to total vector instructions is the workload average vector length. The 4Q93 value is about 68 whereas the C-90 hardware vector length is 128. The measured length exceeds the previous quarter's value by 15%. User programs with vector lengths closer to the hardware length can better exploit the vector performance.

Vector memory load (read) rates are twice as large as vector memory store (write) rates. A FLOP requires, on the average, one memory reference, but it is more likely to be a load than a store. The C-90 architecture provides each CPU with two double-width memory paths for loading data from memory and one memory path for storage; the architecture reserves the fourth memory path for I/O and instruction buffer transfers. The C-90 provides a maximum memory bandwidth of 6 references per CP per CPU. Since the maximum CPU computational rate is 4 floating point operations per CP, the Cray design attempts to ensure CPU-intensive codes will not experience memory-starvation.

The current workload requires a maximum CPU memory bandwidth of 1.3 references per CP. This bandwidth is the maximum workload value for a group of 88 24-hour mea-

surement periods. Some NAS applications require 2.6 references per CP per CPU to maintain their performance rate.

Derived Data:

The table lists several quantities obtained through calculations with the counter data.

Table 5: NAS C-90 4Q93 Daily Average HPM Measurements-Derived Data

Measurement	Unit	Avg	STD	COV	Min	Max
System Availability	Percent	88.70	6.10	6.90	62.00	97.00
System MFLOPS	M/sec	3626.603	375.808	0.104	2184.940	4566.590
Vector Operation Fraction	Percent	91.830	0.813	0.009	90.050	94.020
Scalar Operation Fraction	Percent	8.170	0.813	0.099	5.980	9.950
Vector Operation Rate	M/sec	548.574	36.128	0.066	464.730	656.610
Scalar Operation Rate	M/sec	48.520	2.488	0.051	40.030	54.360
Total Operation Rate	M/sec	597.094	34.219	0.057	516.110	699.900
Instruction Issue Fraction	Percent	23.606	1.041	0.044	20.184	25.823
Hold Issue Fraction	Percent	67.729	1.531	0.023	64.286	72.684
Null Instruction Fraction	Percent	8.764	0.587	0.067	6.926	9.972

Availability is the fraction of time the C-90 operated in user mode. During other times, the C-90 was either idle or executing system calls. The upgrade of the UNICOS operating system from version 7.C.2 to version 8.0 occurred in early December. This upgrade featured a fully multithreaded kernel to reduce the periodic episodes of high system time observed for the 16-CPU C-90 in 3Q93. Examination of system logs indicated no consistent basis for the measured increase in overhead. The probable cause was kernel contention, i.e., the updating of kernel data tables by a single CPU which prevented other CPUs from accessing those same tables and which led to stalling of the CPUs.The Cray System Activity Report (SAR) recorded an average of 11.5% idle for the 8 weeks prior to the UNICOS upgrade and 4.5% following the upgrade to version 8.0. The large COV for availability reflects the problem with UNICOS 7.C.2.

System MFLOPS denotes the system throughput. This rate is the product:

System MFLOPS = MFLOPS/CPU *CPUs*Availability.

The table shows the throughput rate to be 3626 MFLOPS or 23.6% of peak.

The table indicates that 91.8% of the operations were performed in vector mode, an increase of 1% over the previous quarter. The total operation rate of 597 MOPS represents an increase of 4.5% over the previous quarter's rate. Since this rate is about 2.48 OPS/CP, the instruction processor is able to overlap operations despite the large number of hold issue CPs discussed under Table 1.

Of the 597 MOPS, 41% were memory operations. If typical machine operations followed the "Load,Load,Operate, and Store" pattern, 75% of the operations would be memory operations. The reduced amount of memory usage confirms that the compiler is successfully using registers to limit memory operations

A complete accounting of all CPs accumulated by the C-90 CPU while in user mode includes the time spent issuing instructions, the time spent holding instruction issue, and the time spent preparing for the next instruction. Cray terms the latter quantity NIP (Next Instruction Parcel) time and includes in it the CPs spent jumping across instruction buffers, CPs spent fetching words from memory to load the buffers, and CPs spent processing instruction of more than one word in length. Partition of total CP time into these categories can illustrate the reasons for performance differences between two similar workloads. In 4Q93, the C-90 spent about 24% of the user time issuing instructions, 68% of the user time holding issue, and 9% of the user time preparing for the next instruction. This breakdown is consistent with a sequential instruction issue punctuated by delay periods due to resource reservation. The small amount of workload NIP time indicates relatively well-written code.

The 1 GW upgrade in the C-90 memory contributed to the 13% increase in the 4Q93 vector length and a smaller 5% increase in 4Q93 CPU MFLOP rate. The following figure shows the cumulative fraction of user CPU time as a function of memory size. The figure shows that large memory jobs tend to occupy a progressively larger fraction of the CPU time in the months following the memory upgrade.

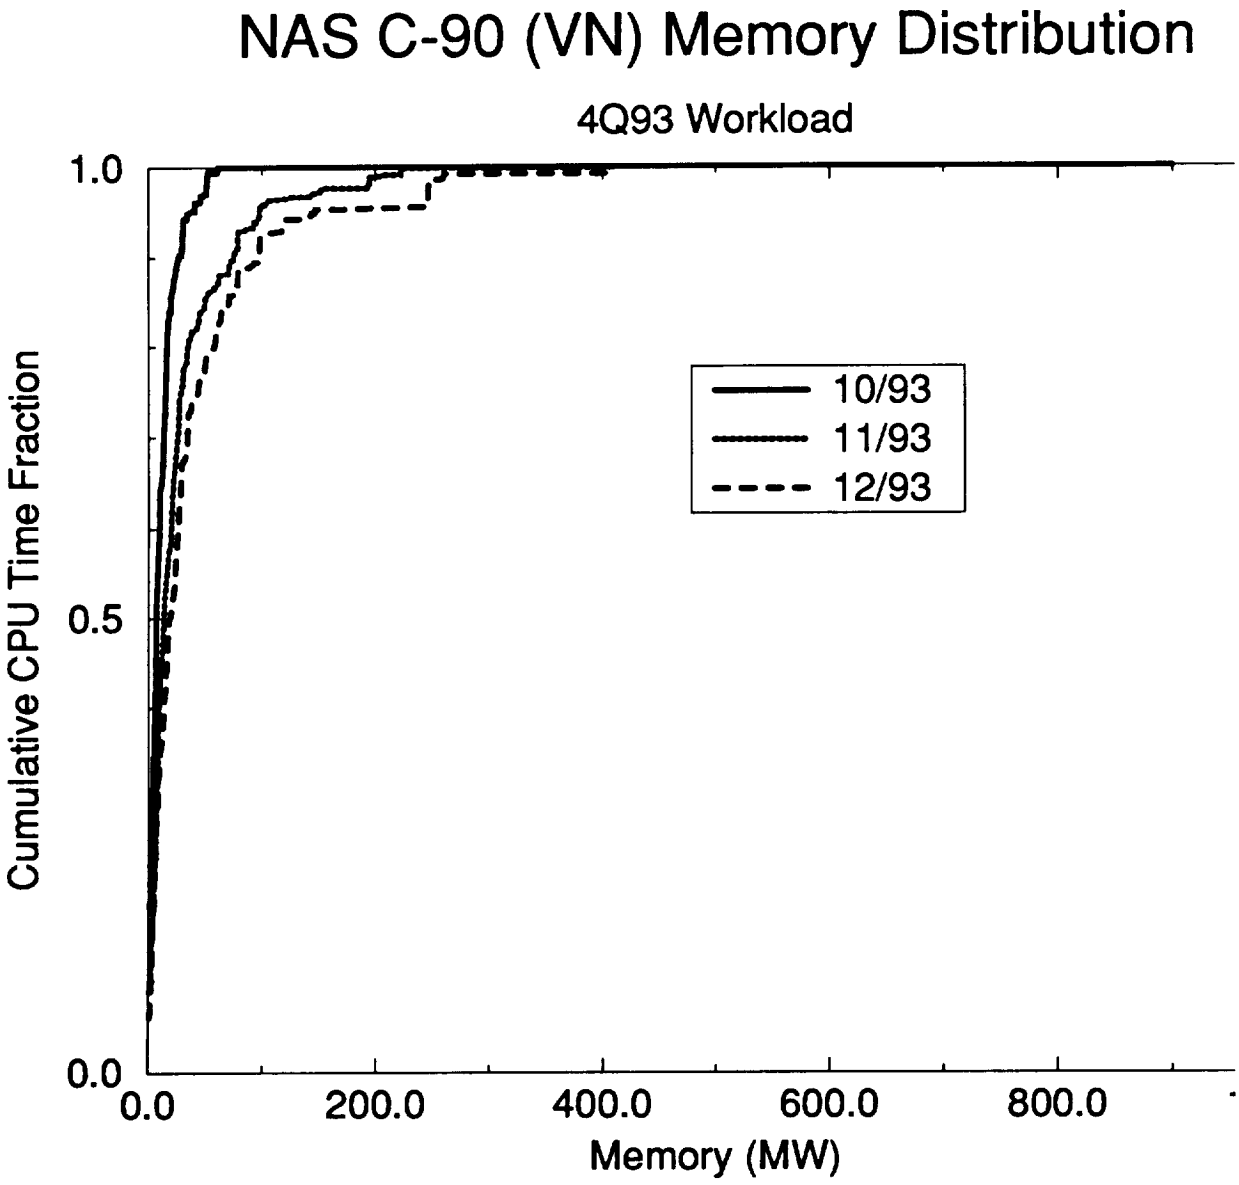
NAS C-90 (VN) Memory Distribution
4Q93 Workload
Cumulative CPU Time Fraction
1.0
0.5
0.0
0.0
200.0
400.0
600.0
800.0
Memory (MW)
10/93
11/93
12/93

3.0 Y-MP Counter Data

The NAS Y-MP (SN 1030) CPUs have a clock period (CP) of 6.000 nanoseconds and the CPUs have a peak speed of 333 MFLOPS. Only one group of counters can be active at any one time and the group chosen for the measurements is the global counter set, Group 0. These counters provide total counts of instructions, operations and memory references.

Group 0:

Table 6 provides 4Q NAS Y-MP per-CPU counter data for Group 0. The table also provides values calculated from the counter data. This architecture has 64-element vector registers and single-width 64-element functional units.

Table 6: NAS Y-MP 4Q93 Daily Average Hardware Measurements-Group 0

Measurement	Unit	Avg	STD	COV	Min	Max
CPU time	Sec	76321.	9329.	0.122	44428.	83484.
Instruction Issue	M/sec	39.119	1.661	0.042	35.990	42.400
Average clock periods/inst	----	4.268	0.178	0.042	3.930	4.630
CP holding issue	Percent	66.514	1.599	0.024	63.360	69.770
Instruction buffer fetches	M/sec	0.313	0.023	0.072	0.270	0.360
Floating Pt. adds	M/sec	43.064	2.741	0.064	38.310	50.000
Floating Pt. multiplies	M/sec	45.433	3.105	0.068	41.000	52.540
Floating Pt. reciprocals	M/sec	2.647	0.369	0.139	1.850	3.280
Floating Pt. Ops/CPU	M/sec	91.144	5.505	0.060	82.210	104.340
CPU memory references	M/sec	92.603	5.546	0.060	83.220	107.980
I/O memory references	M/sec	0.656	0.435	0.664	0.200	2.490

During 4Q93, each of the Y-MP's 8 CPUs performed at an average rate of 91 MFLOPS, a decrease of about 7% over the previous quarter. This decrease is rather substantial for the 30-day Y-MP measurement period and probably reflects the transition of most production codes to the C-90.

The CPU MFLOP rate slightly exceeds the memory reference rate, indicating that each floating point operation requires slightly more than one memory reference. The ratio implies register reuse to avoid main memory accesses.

The I/O memory reference rate of 0.656 Mwords/sec per CPU represents a 33% increase over the last quarter.

Reasons for the large COVs for CPU time and I/O memory reference rate are the same as those provided for the C-90 in Section 2.

Derived Data:

Table 7 lists several quantities obtained through calculations with the counter data.

Table 7: NAS Y-MP 4Q93 Daily Average HPM Measurements-Derived Data

Measurement	Unit	Avg	STD	COV	Min	Max
System Availability	Percent	93.00	2.00	2.10	88.00	97.00
System MFLOPS	M/sec	678.494	43.192	0.064	609.850	804.050
Vector Operation Fraction	Percent	86.				
Scalar Operation Fraction	Percent	14.				
Vector Operation Rate	M/sec	223.				
Scalar Operation Rate	M/sec	39.				
Total Operation Rate	M/sec	262.				
Instruction Issue Fraction	Percent	23.471	0.097	0.042	21.594	25.440
Hold Issue Fraction	Percent	66.514	1.599	0.024	63.360	69.770
Null Instruction Fraction	Percent	10.015	0.667	0.067	8.636	11.348

System MFLOPS denotes the system throughput. This rate is the product:

System MFLOPS = MFLOPS/CPU *CPUs*Availability.

Availability is the fraction of elapsed time in which the Y-MP operated in user mode. The table shows the throughput rate to be 678 MFLOPS or 25.4% of the theoretical peak rate.

The table provides a value of 86% for the fraction of Y-MP operations performed in vector mode based on previous NAS measurements [2].

The total operation rate includes the operations produced by the vector and the scalar instructions. The Y-MP Group 0 output does not provide vector and scalar integer and logical operations. The vector value in Table 7 is the sum of the Group 0 memory operation rate, floating point operation rate, and instruction issue rate. Experience indicates that the latter can serve as a rough approximation to the logical and integer operations not included in Group 0. The aggregate rate of 235 MOPS per CPU represents about 1.40 OPS/CP, indicating that the instruction processor is able to overlap operations despite the large number of hold issue CPs. Of the 234 MOPS, 43% were memory operations. If typical machine operations followed the "Load Operate, and Store" pattern, 75% of the operations would be memory operations. The reduced amount of memory usage confirms that the compiler is successfully using registers to limit memory operations.

As with the C-90, a complete accounting of all CPs accumulated by the CPU while in user mode includes the time spent issuing instructions, the time spent holding instruction issue, and the time spent preparing for the next instruction. The Y-MP spent about 24% of the user time issuing instructions, 68% of the user time holding issue, and 8% of the time preparing for the next instruction. The Null Instruction Fraction, while still constituting a small amount of user CPU time is about 10% larger than that of the C-90.

4.0 Discussion

The 4Q C-90 average CPU performance was about 5% larger than that of the previous two quarters. The following table summarizes some key results:

Table 8: NAS C-90 Key Hardware Performance Results

Measurement	2Q93	3Q93	4Q93
CPU MFLOPS	244	244	255
Percent Vectorization	91.7	91.0	91.8
Vector Length	62.9	59.7	67.6
System Availability	85.6	88.3	88.7
System GFLOPS	3.315	3.442	3.626
System Efficiency	21.6	22.4	23.6

The table shows an increases in CPU MFLOPS and System GFLOPS during 4Q93. At a 90% confidence level, the 4Q confidence intervals for these two quantities lie outside these quantities' confidence intervals for the previous two quarters. Thus, the 4Q performance increases are statistically significant and deserve an explanation.

The table also indicates a substantial (and statistically significant) increase in vector length.Long vector lengths help to amortize the overhead of starting the vector functional

units. The relationship between workload CPU performance and workload vector length contains considerable scatter, but a least squares fit to the data suggests that the increase in vector length can account for most of the CPU performance increase.The introduction of the 1 GW memory assisted in increasing C-90 performance as job memory requirements increased in each of the months comprising 4Q93.

The following table lists the corresponding results for the more mature NAS Y-MP workload.

Table 9: NAS Y-MP Key Hardware Performance Results

Measurement	2Q93	3Q93	4Q93
CPU MFLOPS	97	98	93
Percent Vectorization	86(est.)	86(est.)	86(est.)
Vector Length	46(est.)	46(est.)	46(est.)
System Availability	86	90	93
System GFLOPS	0.665	0.705	0.678
System Efficiency	25.0	26.5	25.4

Since the Y-MP workload shows no trend toward increasing system performance despite an increase in system availability, the vector fraction and vector length, presented in the table as values obtained from 1989 measurements, are probably incorrect.

Although the users are employing similar codes on both machines, the C-90 user programs display vector lengths which are short relative to the longer C-90 hardware length. These programs do not fully exploit the vector hardware and lead to a lower efficiency relative to the Y-MP, 23.6% vs. 25.4%.

Most NAS C-90 users (the exceptions are the well-written Langley codes) can increase the vector lengths of their Cray codes by overindexing DO loops. This term describes the reduction of a nested loop into a single loop. As an illustrative example, consider

```
      DIMENSION S(67,67),X(67,67),Y(67,67)
      DO 100 I=1,67
      DO 100 J=1,67
        S(I,J)=A*X(I,J)+Y(I,J)
  100 CONTINUE
```

With compiler optimization off, this loop has a vector length of 67. With overindexing, the above loop would appear as

```
      DIMENSION A(67,67)
      DO 100 IJ=1,67*67
        S(IJ)=A*X(IJ)+Y(IJ)
  100 CONTINUE
```

On the C-90, this loop would have vector length of 123 and would display a performance increase factor of 1.2.

Measured NAS performance data indicate that the current workload has the potential to achieve a performance increase of 1.3 due to overindexing. The drawback is that overindexing is a nonstandard Fortran construct and portability to other platforms is not guaranteed.

NAS can employ multiple approaches to promote increased workload vector length. Instructions concerning the use of overindexing should be included in the NAS User Guide and in the NAS News. NAS can modify the HPM performance summary provided by Cray to report short vector length programs by user and also by group. Consultants could target guidance to users and groups with deficient vector lengths.

5.0 Conclusion

The NAS C-90 workload displayed a 5% increase in throughput during 4Q93. During this period, workload vector length increased by 13% and workload vector fraction increased by 1% over 3Q measurements. Analysis indicates the performance increase was statistically significant and was due to the increase in vector length. The increased vector length coincided with the introduction of 1 GW memory. Despite this increase, insufficient vector length continues to hinder C-90 performance and NAS should employ several approaches to remedy this problem. Analysis of memory-related delays in issuing instructions and fetching data indicated that the memory is not a bottleneck for the current NAS workload.

6.0 Acknowledgment

Thanks to David Barkai and Duane Carbon for reviewing this paper.

7.0 References

[1] Cray Research, Inc. UNICOS Performance Utilities Reference Manual, SR-2040 7.0, 1992.

[2] Bergeron, R.J."Performance Analysis of the NAS Y-MP Workload", NASA Ames Research Center, NAS Systems Division Report RND-90-009, December, 1990.

RND TECHNICAL REPORT

Title: The Performance of the NAS High Speed Processors in 4Q93

Author:
Robert J. Bergeron

Reviewers:
"I have carefully and thoroughly reviewed this technical report. I have worked with the author(s) to ensure clarity of presentation and technical accuracy. I take personal responsibility for the quality of this document."

Signed: David Barkai

Name: David Barkai

Signed: Duane P. Carbon

Name: Duane Carbon

Branch Chief:

Approved: Bruce Tuchler

Date & TR Number:

July 1994 RND-94-012